arXiv:1410.7990v1 [cs.DB] 27 Oct 2014

Linked Data Integration with Conflicts

Jan Michelfeit*, Tomáš Knap, Martin Nečaský

XML and Web Engineering Research Group, Charles University in Prague, Malostranské Náměstí 25, 118 00 Praha 1, Czech Republic

Abstract

Linked Data have emerged as a successful publication format and one of its main strengths is its fitness for integration of data from multiple sources. This gives them a great potential both for semantic applications and the enterprise environment where data integration is crucial. Linked Data integration poses new challenges, however, and new algorithms and tools covering all steps of the integration process need to be developed. This paper explores Linked Data integration and its specifics. We focus on data fusion and conflict resolution: two novel algorithms for Linked Data fusion with provenance tracking and quality assessment of fused data are proposed. The algorithms are implemented as part of the ODCleanStore framework and evaluated on real Linked Open Data.

Keywords: Linked Data, data integration, conflict resolution, data quality, data fusion

1. Introduction

More and more valuable datasets are being published on the Web and often their usefulness increases dramatically when data can be combined from multiple sources. Similarly, data in the enterprise environment are often distributed across many independent systems and their full value and potential can be exploited only when integrated together. Linked Data are an instrument specifically designed to facilitate linking and data integration across datasets and provide many advantages for data integration over alternative approaches.

Data integration in relational databases is a well explored field with mature tools and frameworks covering all steps of integration. On the other hand, Linked Data integration has still open challenges. One is the resolution of conflicts and uncertainties emerging during integration. Another challenge is information quality whose importance is significant especially in the open Web environment. Both data consumer and conflict resolution tools need support in decisions about which data are worth using. We must also face technical challenges – integration must be efficient and a pay-as-you-go approach [21] may be essential for adoption in practice.

This paper aims to fill in the missing pieces in the data fusion step of Linked Data integration, and address the outlined challenges. We provide both a practical tool and a theoretical framework as a basis for further research. The main contributions are:

1. An algorithm realizing the data fusion step in Linked Data integration with resolution of conflicts and provenance tracking. The algorithm deals with usage of different identifiers and schemata, and conflicting or missing values in data sources.
2. A conflict-based quality assessment algorithm which leverages information available during data fusion. We introduce the concept of *F-quality* as a measure of quality of fused data as opposed to quality of source data.
3. Overview of the specifics of Linked Data integration in comparison with relational databases.

We show how the proposed algorithms work together to improve both the abilities of conflict resolution and quality assessment. Both algorithms are implemented as part of a Linked Data integration framework ODCleanStore and evaluated on real Linked Open Data.

The paper is structured as follows: Section 2 examines the data integration and data fusion process for Linked Data and introduces the ODCleanStore framework. Section 3 gives an overview of related work. Section 4 describes in detail the proposed data fusion algorithm and Section 5 covers conflict-based quality as-

*Corresponding author. Tel: +420 732 737 415

Email addresses: michelfeit@ksi.mff.cuni.cz (Jan Michelfeit), knap@ksi.mff.cuni.cz (Tomáš Knap), necasky@ksi.mff.cuni.cz (Martin Nečaský)

sessment. Section 6 presents experimental results. Section 7 compares the specifics of Linked Data integration to relational databases and we summarize our results in Section 8.

2. Data Integration & ODCleanStore Framework

Data integration is about combining data from different sources to a unified view. The main challenges that must be addressed in order to achieve a unified integrated view include:

- Technical and semantical heterogeneity of data.
- Schema, identity, and data conflicts.
- Incorrect or otherwise flawed data.
- Identification of a target schema and schema translation.
- Presentation of results.

Data integration systems use various combinations of steps to cope with these challenges. Generally, these steps are:

1. Schema mapping.
2. Data source selection and data retrieval.
3. Data transformations and schema translation.
4. Duplicate detection (object identity resolution).
5. Quality assessment.
6. Data fusion and conflict resolution.
7. Result loading or visualization.

We aim to create a Linked Data integration framework covering all these steps in *ODCleanStore* [15]. ODCleanStore is a server application for integration and management of Linked Data. It accepts Linked Data as RDF,[1] processes them in a customizable pipeline of data processing units and saves the result to a data store. Users are provided with integrated views on the processed data that are generated on demand by the data fusion component presented in this paper. This *partially materialized* approach gives us the flexibility of being able to add or change stored data or even schema mappings at any time without the need to re-run data fusion on the whole data store, thus giving us the advantages of gradual evolution in a pay-as-you-go manner. The whole process is designed so that the trust aspect is supported with provenance tracking and quality assessment of integrated data.

[1] `http://www.w3.org/TR/rdf-syntax/`

2.1. Data Fusion

The contribution of this paper covers the data fusion phase with conflict resolution and a conflict-aware quality assessment of fused data. We present new algorithms that are implemented in ODCleanStore and are also available as a standalone tool *ODCS-FusionTool*.[2]

Data fusion is the step where actual data merging happens – multiple records representing the same real-world object are combined into a single, consistent, and clean representation [3]. In order to fulfill this definition, we need to establish a representation of a record, purge uncertain or low-quality values, and resolve identity and other conflicts. Therefore we regard conflict resolution as a subtask of data fusion.

Conflicts in data emerge during the data fusion phase and can be classified as schema, identity, and data conflicts. Schema conflicts are caused by different source data schemata – different attribute names, data representations (e.g., one or two attributes for name and surname), or semantics (e.g., units). Identity conflicts are a result of different identifiers used for the same real-world objects. Finally, data conflicts occur when different conflicting values exist for an attribute of one object.

Conflict can be resolved on *entity* or *attribute* level by a *resolution function*. Resolution functions can be classified as *deciding functions*, which can only choose values from the input such as the maximum value, or *mediating functions*, which may produce new values such as average or sum [3].

The basic structure of the presented data fusion algorithm and its inputs are outlined in Figure 1. In the context of the ODCleanStore framework (Figure 1a), data fusion is executed at query time before query results are returned to the user. Application of the algorithm can be facilitated by preprocessing steps such as data transformations or schema translation. The standalone implementation (Figure 1b) can be used for batch processing on raw RDF data stored in SPARQL endpoints or files. These can optionally also contain the metadata and mappings to be used in conflict resolution.

2.2. Data Quality

Measurement of information quality plays a fundamental role in data integration. Quality score may be the factor deciding which value to pick during conflict resolution [19]; even if not quantified by a scoring function, quality indicators and metadata are commonly used in

[2] For the rest of this paper, we will use the name *ODCS-FusionTool* to refer to both the implementation of data fusion in ODCleanStore and in ODCS-FusionTool.

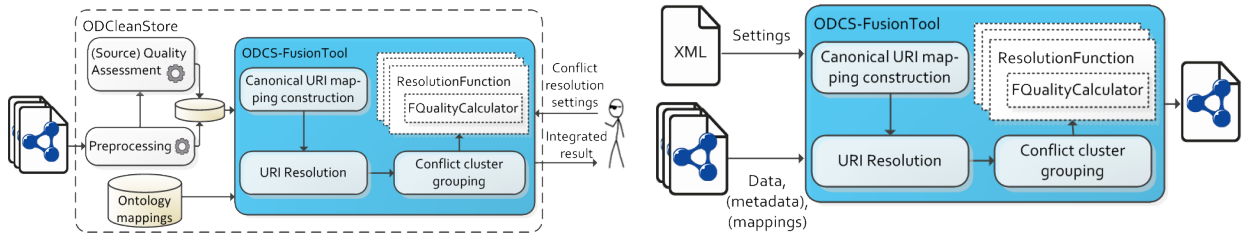

(a) Data fusion in the ODCleanStore framework

(b) Data fusion with the standalone tool ODCS-FusionTool

Figure 1: High-level overview of the ODCS-FusionTool implementation in the context of ODCleanStore and when used independently.
ODCS-FusionTool constructs a mapping to canonical URIs from input same-as links and ontology mappings in the first step which is used to resolve identifier and schema conflicts in the second step. The result is grouped to clusters of possibly conflicting triples where conflicts are resolved by a resolution function according to configuration. Resolution function also produces provenance and quality information about the integrated result.

conflict resolution [1, 2, 7, 28]. These indicators may be both assessed automatically or entered manually by an expert user.

The quality of data produced *from* a data integration process is also very important. It helps data consumers in decisions whether to trust a particular piece of information, or for overall quality assessment before further data processing.

We introduce the concept of *F-quality* as a measure of quality of integrated and fused data. We also propose a method of F-quality assessment which has been implemented and evaluated in conjunction with our data fusion algorithm in ODCS-FusionTool. The assessment method is based on the conflict resolution context represented by conflicting statements, quality-related metadata, and user settings given to a conflict resolution function. Compared to other quality assessment methods, leveraging the conflict resolution context allows data fusion to make decision even in the absence quality-related metadata using a voting-like approach.

3. Related Work

3.1. Linked Data Integration

In the young area of Linked Data, tools focus on individual aspects of data integration, e.g., duplicate detection or federated querying, but apart from ODCleanStore only LDIF [26] offers a complex data integration solution. The same goes for data fusion and conflict resolution – to the best of our knowledge, only Sieve [19], a part of LDIF, solves this task.

Linked Data Ingteration Framework (LDIF) is a framework producing homogenized views over heterogeneous data from diverse sources with a modular architecture. Modules include data retrieval, schema translation, duplicate detection, quality assessment, and data fusion with conflict resolution (Sieve).

Sieve consists of two parts that run in sequence. First, Quality Assessment produces quality scores from user-selected metadata and configuration. The score calculation is based on concepts such as *assessment metric*, *data quality indicator*, or *scoring function*. In the second step, Data Fusion Module uses the quality scores in order to perform user-configurable conflict resolution. Sieve offers an extensible set of *fusion functions* used for resolution of conflicts.

ODCS-FusionTool presented in this paper has a similar purpose as Sieve – data fusion with conflict resolution. Due to a different intended use of the two, ODCS-FusionTool has several unique features compared to Sieve. Sieve is used to fuse data as part of an ETL process realized by LDIF, i.e. conflicts are resolved when data are loaded into a data store. ODCS-FusionTool, on the other hand, is designed as an independent library which can be used either when loading data, or at query time, supporting a pay-as-you-go approach [21].

In Sieve, Quality Assessment produces scores through user-configured scoring functions and then Data Fusion is able to use them during conflict resolution. ODCS-FusionTool doesn't include quality assessment (it is a separate module in ODCleanStore) but uses arbitrary given RDF metadata as its input. More importantly, it produces quality scores of the *fused* data, as opposed to source data. This way we can leverage the context information available in the fusion phase and reflect it in quality calculation, e.g. the quality may be different for a value confirmed by several sources and a value that

is contradicted by other sources. In addition, it enables quality-based decisions even when source metadata are missing, as opposed to Sieve.

LDIF optionally outputs provenance metadata and results of Quality Assessment of source data. This is beyond the scope of ODCS-FusionTool (metadata are given on demand in ODCleanStore) but it outputs provenance of *resolved* data so that data consumers can track where exactly the fused quads originated from. Another difference is that ODCS-FusionTool can track quality to the quad level whereas LDIF tracks only to named graphs.

3.2. Conflict Resolution

Sieve, described in Section 3.1, is the only project that addresses conflict resolution for Linked Data. This section lists non-Linked Data integration systems that acknowledge conflicts and provide means for their resolution, sorted by descending year of publication.

Fusionplex [20] integrates heterogeneous information sources with resolution of factual inconsistencies. It uses metadata such as timeliness, accuracy or cost to compute *data utility*. Data can be filtered based on data utility or the actual values and fused by a selected fusion function (average, maximum, any, ...). Data inconsistencies are handled in several steps: (1) Tuples sharing the same key are grouped to *polytupes*. (2) Each polytupe is purged according to user preferences. (3) Values for each attribute are purged and fused to a single value, leaving a single tuple per polytupe.

The Humboldt Merger [1] (HumMer) allows ad-hoc, declarative fusion of conflicting data using an extension to SQL. It features schema matching, duplicate detection, and data fusion/conflict resolution steps. All steps are executed ad-hoc at runtime in response to a user query. The user query is expressed in SQL extended with a `FUSE BY` statement and conflict resolution specifications. HumMer provides a range of conflict resolution functions including source preference, voting, most recent, or standard SQL aggregations. HumMer goes one step further then other systems and optionally visualizes each intermediate step of data fusion with the possibility to interfere.

DaQuinCIS [23] has a module for resolution of data conflicts based on a custom D^2Q data quality model.

FraQL [22] is a lightweight SQL extension for defining integrated object-relational schemata as well as formulating queries over them. FraQL recognizes several kinds of conflicts including both schema and data conflicts. Schema level conflicts can be resolved declaratively by attribute mappings or procedurally, data conflicts by user-defined reconciliation functions.

Mediation framework **AURORA** [33] proposes extension of SQL to support *conflict tolerant queries*, which produce conflict-free results in the presence of conflicts. User specifies a fusion function for columns in the query result (e.g., AVG, MAX, ANY, DISCARD, user-defined). In addition, a tolerance strategy (one of *HighConfidence*, *RandomEvidence*, *PosibleAtAll*) may be given for evaluation of query conditions.

OO_{RA} [17] is an extended object-oriented data model and SQL extension to handle attribute-level conflicts where attribute resolution functions can be given.

Hermes [28] integrates information from diverse sources (not limited to databases) employing declarative rule-based language for defining mediators in an extensible way with the ability to define conflict resolution rules. The mediator may use predefined or custom strategies (e.g., LATEST, MAX).

TSIMMIS [7] uses a custom *Object Exchange Model* based on quadruples (label, type, value, object-iid) extracted from the underlying sources. It recognizes possible inter-source duplicates and value conflicts, although only simple source preference is supported.

In **Multibase** [8], data describing the same type of entities are integrated using a principle of generalization (global classes generalize local ones) and SQL aggregation functions (SUM, AVG, ...) for inter-source conflicts.

ODCS-FusionTool is a more mature system in its modular architecture where quality computation and fusion functions are independent and pluggable. It also leverages OWL[3] to resolve schema and identity conflicts based on mappings. Finally, it is novel in that the quality calculation can leverage the context of multiple data sources and it produces quality and provenance of the actual fused data rather then just source data. Several projects use declarative configuration by extending SQL. Extending SPARQL[4] for such purpose is a topic for further research. Networked graphs [25] are a mechanism which can be used to declaratively integrate RDF data using SPARQL, although it doesn't support data conflict resolution directly.

3.3. Conflict-Based Quality

A general framework for data fusion by Ronald Yager [32] proposes a voting-like process to determine the best values. A total *support* is computed for each solution taking conflicts with other values into consideration. Using the approach described in Section 5.2,

[3] `http://www.w3.org/TR/owl-features/`
[4] `http://www.w3.org/TR/rdf-sparql-query/`

we have independently come to a solution that overlaps with Yager's framework.

In Yager's framework, a fusion engine has multiple sources on its input, each claiming a value a_i, and chooses the best fused solution a based on user requirements, source credibilities, a proximity knowledge base and a knowledge of reasonableness. A *support* of each source is computed for each possible a. The support is based on proximity of a and a_i values weighted by credibility of sources. Support values are then combined together to produce a total support which may further be combined with a reasonableness value for each a_i.

The results produced from F-quality assessment proposed in this paper can be regarded as a special case of Yager's *support* with the following differences: [32] includes a domain-specific reasonableness in the calculation; F-quality assessment instead goes one step further in considering confirmation of values by multiple sources. [32] analyses requirements on operators used in the support function and its properties. These results apply to F-quality as well.

4. Data Fusion and Conflict Resolution Algorithm

Let us recapitulate the main challenges of Linked Data fusion that need to be addressed:

1. Different identifying URIs are used to represent the same real-world entities.
2. Different schemata are used to describe data.
3. Data conflicts emerge when RDF triples sharing the same subject and predicate have inconsistent values in place of the object.

Context that the fusion algorithm can leverage consists of the data to be integrated, metadata and mappings between resource URIs and property URIs from the used schemata, and conflict resolution settings. In this section, we introduce the necessary terminology, describe our proposed algorithm and analyze its time & memory complexity.

We will demonstrate the algorithm on a simple example used throughout this section. Let us consider Example 1 with five statements[5] from three sources that represent a label and the geographical longitude of the city of Berlin. We can see different URIs are used to identify Berlin, properties from different vocabularies are used, and two conflicting values for each property are provided. Our task will be to fuse data about Berlin and provide a single best value for each property of Berlin, i.e. label and longitude.

[5] We use the TriG notation (`http://www.w3.org/TR/trig/`)

Example 1: Sample data before data fusion

```
GRAPH <\protect\vrule width0pt\protect\href{http://dbpedia.or
  db:Berlin  rdfs:label  "Berlin".
  db:Berlin  geo:long    "13.399". }
GRAPH <\protect\vrule width0pt\protect\href{http://rdf.freeba
  fb:en.berlin  rdfs:label  "Berlin".
  fb:en.berlin  fbgeo:long  "13.383". }
GRAPH <\protect\vrule width0pt\protect\href{http://data.nytim
  nyt:N50987  skos:label  "Berlin (Germany)" }
```

4.1. Formalism

Definition 1 (RDF nodes). Let U, B and L be sets of all URI references, blank nodes and RDF literals, respectively. Sets U, B and L are pairwise disjoint. An RDF node is an element of their union $N = U \cup B \cup L$.

Definition 2 (RDF triple). An RDF triple is a statement expressing that a resource has a property with a certain value. Formally, the set of all triples is

$$Triples = (U \cup B) \times U \times (U \cup B \cup L).$$

We denote the elements of each triple as *subject* (resource of interest), *predicate* (property), and *object* (value of the property), respectively.

Definition 3 (RDF graph, Named Graph). A subset G of $Triples$ can be represented as a directed labeled graph and we refer to it as an *RDF graph*.

A *Named Graph* is a pair (G, n) where $G \subset Triples$ is an RDF graph and $n \in U$. We say that graph G is *named* n. Named Graphs cannot share blank nodes, i.e. blank nodes in triples from (G_1, n_1) are distinct from those in (G_2, n_2) for $n_1 \neq n_2$.

Definition 4 (Quad). Let *quad* denote a quadruple (s, p, o, g) such that there is a triple (s, p, o) in RDF graph G named g. The set of all quads is denoted:

$$Quads = (U \cup B) \times U \times (U \cup B \cup L) \times U$$

Definition 5. Let G be an RDF graph. We define functions $subjects(G)$, $predicates(G)$ and $objects(G)$ by the following formulas: $subjects(G) = \{s \mid (s, p, o) \in G\}$, $predicates(G) = \{p \mid (s, p, o) \in G\}$, $objects(G) = \{o \mid (s, p, o) \in G\}$ We define these functions on a set of quads $Q \subset Quads$ analogously and $graphs(Q) = \{g \mid (s, p, o, g) \in Q\}$.

The result of the conflict resolution algorithm implemented in ODCS-FusionTool contains not only the RDF triples with conflicts resolved according to the given conflict resolution strategy but also:

1. Names of the Named Graphs each triple was selected from or derived from.
2. A quality value for each triple based on conflicting values and input metadata; quality is expressed as a value from an ordered space $C = [0; 1]$. This quality (formalized in Definition 9) expresses how trustworthy the respective triple is with regards to other conflicting values and input metadata (such as data source quality).

In order to convey this information in the result, we introduce *resolved quads* as output of the conflict resolution algorithm.

Definition 6 (Resolved quad). *Resolved quad* is a triple (q, S, c) from the space of all result quads denoted $ResolvedQuads$ defined as

$$ResolvedQuads = Quads \times \mathcal{P}(U) \times C,$$

where q is a result quad, S is the set of names of Named Graphs q was selected or derived from ($\mathcal{P}$ denotes the power set) and c is the quality.

Example 2 shows how a resolved quads produced for data from Example 1 may look like.

Example 2: Resolved quad

```
((dbpedia:Berlin, rdfs:label, "Berlin", ex:1),
 {http://dbpedia.org,\protect\vrule width0pt\protect\href{http://rdf.freebase.com}{http://rdf.freebase.com}},
 0.71)
((dbpedia:Berlin, geo:long, 13.391, ex:2),
 {http://dbpedia.org,\protect\vrule width0pt\protect\href{http://rdf.freebase.com}{http://rdf.freebase.com}},
 0.85)
```

The proposed algorithm deals with conflicting property values, i.e. conflicts in place of quad objects. This is sufficient to examine all triples and it naturally corresponds to attribute value resolution in traditional conflict resolution [3]. Therefore quads which are in conflict must share the same subject and predicate. This motivates the following definition.

Definition 7 (Object conflict cluster). *Object conflict cluster* in a set of quads $Q \subseteq Quads$ is a maximal subset $CC \subseteq Q$ such that $|subjects(CC)| = 1$ and $|predicates(CC)| = 1$.

In other words, all quads in a conflict cluster share the same subject and predicate. We will denote the conflict cluster with subject s and predicate p as $CC_{s,p}$.

Conflict resolution functions in the context of relational databases typically operate on attribute values from the attribute's domain or on a set of tuples representing a database record. For Linked Data, the former approach would lead to loss of provenance information. Therefore whole quads must be given to a conflict resolution function rather than conflicting (object) values.

Conflict resolution functions also need additional metadata and context information in order to have enough expressive power to implement all desired resolution strategies. These metadata and context data can also be modeled in RDF as quads. Therefore, we use the following definition of a conflict resolution function.

Definition 8 (Conflict resolution function). *Conflict resolution function* f is a function

$$f : \mathcal{P}(Quads) \times \mathcal{P}(Quads) \to \mathcal{P}(ResolvedQuads).$$

Object conflict resolution function f' is a partial function

$$f' : \mathcal{P}(Quads) \times \mathcal{P}(Quads) \to \mathcal{P}(ResolvedQuads)$$

such that the following holds for $f'(CC, M)$:

1. If CC is an object conflict cluster with subject s and predicate p, then $f'(CC, M)$ is also a conflict cluster with the same subject and predicate.
2. If $CC \neq \emptyset$ and CC is not an object conflict cluster in $Quads$, then $f'(CC, M)$ is undefined.

The first argument CC of an (object) conflict resolution function represents the (possibly) conflicting quads to be resolved. The second argument M represents metadata and context information. The distinction between M and CC is necessary as quads in M may be necessary to decide about resolution of quads in CC but they are not meant to be resolved themselves.

Because our algorithm considers only object conflicts, we further refer to object conflict resolution functions simply as *resolution functions* and to object conflict clusters as *conflict clusters* with no ambiguity.

Resolution functions can be classified as *deciding*, which can only choose values from the input such as the maximum value, or *mediating*, which may produce new values such as average or sum [3].

A comprehensive list of resolution functions relevant for Linked Data can be found in Appendix A.

4.2. Input & Output

The presented conflict resolution algorithm takes the following inputs:

1. Collection of *quads* to be resolved.
2. Metadata represented as a collection of *quads*.

3. Mappings between URI resources.
4. Conflict resolution policy which specifies the default conflict resolution strategy and optionally per-property resolution strategies.

Sample of what values can be given for each of these items can be found in Example 3.

Example 3: Sample data fusion algorithm inputs

1. Data: see Example 1
2. Metadata:
```
<http://dbpedia.org> odcs:score "0.9",
<http://rdf.freebase.com> odcs:score "0.8",
<http://data.nytimes.com> odcs:score "0.8"
```
3. Mappings:
```
rdfs:label odcs:equivalent skos:prefLabel,
geo:long   odcs:equivalent fbgeo:long,
db:Berlin owl:sameAs fb:en.berlin,
db:Berlin owl:sameAs nyt:N50987
```
4. Conflict resolution policy:
```
<Resolution function="BEST">
<Property id="rdfs:label"/></Resolution>
<Resolution function="AVG">
<Property id="geo:long"/></Resolution>
```

Metadata can be anything the chosen conflict resolution function needs to produce an appropriate output and assess its quality. For example, it can be the timestamp for each named graph occurring in the input when using the LATEST resolution function, and named graph quality or user preference used in fused data quality assessment (the quality assessment algorithm proposed in Section 5 can work even in the absence of such metadata).

The mappings between URI resources express both results of duplicate detection and schema mappings. They are represented as RDF triples with `owl:sameAs`[6] as their predicate – e.g., a triple (*s*, `owl:sameAs`, *o*) states that *s* and *o* represent the same thing and are equal for the purposes of conflict resolution. Sometimes it is inconvenient to use `owl:sameAs` – e.g., to map `dbp:father` and `dbp:mother` to target property `ex:parent` it is incorrect to state that these two are the same. A special property `odcs:equivalent` can be used in such cases instead.[7]

A *resolution strategy* defines how conflicts shall be resolved. Most importantly, it specifies the *resolution function* to be used.

[6] URI namespace prefixes used in this paper can be resolved with `http://prefix.cc/`

[7] For simplicity, we will further refer to any mapping triples only as *`owl:sameAs` links*.

The chosen resolution function may not be applicable for some values – e.g., function calculating the numeric average is not applicable to string values. The *aggregation error strategy* gives the desired behavior in that case. Inappropriate values may be either discarded or propagated to the output unchanged.

Other parameters affecting, e.g., fused data quality calculation can also be given.

The output of the algorithm is a collection of *resolved quads* (Definition 6), i.e. quads resolved by their respective resolution function together with a quality value and source graph names for each quad.

4.3. High-level Overview

The algorithm includes the following steps in order to deal with the challenges outlined earlier:

1. Replace URI resources connected by a path of `owl:sameAs` links with a single URI (referred to as the *canonical URI*).
2. Remove duplicate identical quads.
3. Group input quads into conflict clusters.
4. For each conflict cluster $CC_{s,p}$:
   (a) Choose a resolution function according to conflict resolution policy, predicate p and `owl:sameAs` mapping.
   (b) Apply the resolution function. This includes resolution of conflicts, calculation of quality and provenance tracking of resolved quads.
   (c) Add the function's output to the result.

Figure 2 shows how the algorithm is applied to our example. Figure 2b shows the result after application of steps 1-3: URIs are replaced with canonical variants and sorted triples are grouped into conflict clusters. The result in 2c is produced by selecting the highest-quality value for label and average for longitude.

4.4. Algorithm Description

Listing 1 is a formal description of the presented conflict resolution algorithm. It uses functions *CanonicalMapping*() to construct canonical URI mappings from input `owl:sameAs` links (Listing 2), *UriResolution*() to resolve identifier conflicts (Listing 3), groups data into conflict clusters, and finally applies an appropriate resolution function to each cluster.

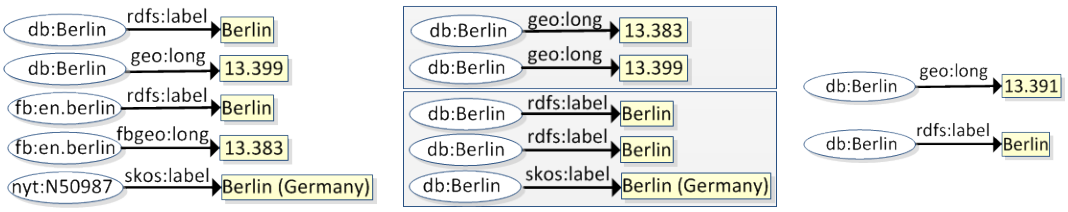

(a) (b) (c)

Figure 2: An example of RDF triples before, during, and after data fusion.

Listing 1: Conflict resolution algorithm

Input: Quads to be resolved Q, metadata quads M, set of `owl:sameAs` links SA, default conflict resolution strategy $defaultStrategy$, partial function $strategy(p)$ returning resolution strategy for predicate p.
Output: Collection of resolved quads.

1: $result \leftarrow \emptyset$
2: $canonical \leftarrow CanonicalMapping(SA)$
3: $Q \leftarrow UriResolution(Q, canonical)$.
4: Sort Q lexicographically by quad subject, predicate, object and graph name.
5: Remove duplicate quads from Q.
6: **for all** s, p such that $(s, p, o, g) \in Q$ **do**
7: $\quad CC_{s,p} \leftarrow \{(s, p, o, g) \mid (s, p, o, g) \in Q\}$
8: $\quad$ **if** $\exists p' : canonical(p') = p$ **and** $strategy(p')$ is defined **then**
9: $\quad\quad strategy \leftarrow strategy(p')$.
10: $\quad$ **else**
11: $\quad\quad strategy \leftarrow defaultStrategy$.
12: $\quad$ **end if**
13: $\quad f \leftarrow$ resolution function according to $strategy$.
14: $\quad result \leftarrow result \cup f(CC_{s,p}, M)$
15: **end for**

Listing 2: Function $CanonicalMapping()$

Input: Set of `owl:sameAs` links SA.
Output: Partial function $canonical : U \rightarrow U$.

1: $canonical \leftarrow \emptyset$.
2: Create graph $H = (V, E)$ with vertices $V = subjects(SA) \cup objects(SA)$, and edges $E = \{\{s, o\} \mid (s, \texttt{owl:sameAs}, o) \in SA\}$.
3: Find set of weakly connected components C in H.
4: **for all** connected component $C = (V_C, E_C) \in C$ **do**
5: $\quad$ Choose URI $c \in V_C$ as the canonical URI for C.
6: $\quad$ **for all** $v \in V_C$ **do**
7: $\quad\quad canonical(v) \leftarrow c$
8: $\quad$ **end for**
9: **end for**
10: **return** $canonical$.

Listing 3: Function $UriResolution()$

Input: Collection of quads to be resolved Q, canonical URI mapping $canonical : U \rightarrow U$.
Output: Collection of quads.

1: **for all** input quad $q = (s, p, o, g) \in Q$ **do**
2: $\quad$ **if** s is a URI resource and $canonical(s)$ is defined **then**
3: $\quad\quad$ Replace s in q with $canonical(s)$.
4: $\quad$ **end if**
5: $\quad$ Repeat steps 2-4 for predicate p and object o.
6: **end for**
7: **return** Q.

Construction of Canonical URI Mapping. The first step of the algorithm is the creation of canonical URI mappings (line 2 in Listing 1). Data from multiple sources may use different URIs to represent the same concepts (identifier conflicts) and different predicate URIs for the same property (schema conflicts). We need to collapse URIs representing the same entity into a single *canonical URI*. Mapping from URI resources to their canonical URI is created from input `owl:sameAs` links by function $CanonicalMapping()$.

The implementation of $CanonicalMapping(SA)$ is based on a disjoint-set data structure which keeps track of elements partitioned into a number of disjoint subsets, with operations *Find* to determine the set an element belongs to, and *Union* to join two subsets into one. A forest-based implementation of disjoint-set can provide $O(\alpha(n))$ amortized time per operation, where α is the inverse Ackermann function [29]. $CanonicalMapping()$ uses such data structure to partition URIs into sets. For each link $(s, \texttt{owl:sameAs}, o)$, $CanonicalMapping()$ calls $Union(s, o)$. Each subset then represents one weakly connected component of the `owl:sameAs` links graph. One member of each set is

chosen as the canonical URI.

URI Resolution. The second step of the algorithm (line 3 in Listing 1, implementation in Listing 3) resolves identifier conflicts and schema conflicts according to given canonical URI mapping and prepares input for the application of a resolution function. This part is independent on the given conflict resolution policy.

Application of Resolution Function. The rest of the conflict resolution algorithm in lines 6-15 of Listing 1 groups quads into conflict clusters and applies a resolution function.

On line 4, quads in Q are sorted lexicographically by their subject, predicate, object and graph name. This allows efficient removal of duplicities on line 5 (which may appear after the URI resolution step), and clustering of conflicting quads on lines 6-7 as quads from one object conflict cluster will be adjacent after sort.

The filtering procedure benefits from already resolved identifier conflicts.

The purpose of lines 8-13 is the selection of an appropriate resolution function. The user who entered resolution strategies for each property may have used a property URI different from what the algorithm selected as a canonical URI. Therefore we locate the appropriate resolution function using *canonical* mapping.

The final step is the application of a resolution function. Implementations of resolution functions can be very diverse (see Appendix A) but with respect to the expected output, they need to (1) produce result quads according to the resolution strategy they implement, (2) calculate quality of each result quad, and (3) keep track of Named Graph each result quad originates from.

Methods how to calculate quality can also be diverse and depend on the resolution function and the task at hand. Our proposed quality assessment method can be found in Section 5.

The architecture of ODCS-FusionTool is flexible and both new resolution functions and quality assessment methods can be plugged in.

4.5. Time & Memory Complexity

In this section, we show that the time complexity of the algorithm is $O(n \log n + l + T(n))$ and memory complexity is input size plus $O(l)$ with few assumptions.

1. Let $n = |Q|$ be the number of quads to be resolved.
2. Let $l = |SA|$ be the number of `owl:sameAs` links.
3. Let $c = |\{(s, p);\ (s, p, o, g) \in Q\}|$ be the number of conflict clusters.
4. Let c_i be the size of i-th conflict cluster.
5. Let $T(x)$ be an upper bound of applied resolution functions' complexity.

We assume the number of defined per-property resolution strategies is constant (independent on n and l), and assume using a disjoint-set data structure with amortized time of operations $\alpha(x)$ as described above.

Time complexity. There are several main operations affecting time complexity. The first is creation of canonical URI mapping in $O(l\,\alpha(l))$; the second is translation to canonical URIs in $O(n\,\alpha(l))$ (at most 3 lookups for each quad); third is sorting of quads in $O(n \log n)$; the last one is application of resolution functions in total time $O\left(\sum_{i=1}^{c} T(c_i)\right)$. Therefore the complexity is

$$O\left(n \log n + (l + n)\,\alpha(l) + \sum_{i=1}^{c} T(c_i)\right).$$

We can simplify this formula with two other simple assumptions. Function $\alpha(l)$ grows very slowly and is less than five for all practical purposes. Let us further suppose that $T(x) = x^n$ (T is polynomial). Using the binomial theorem, we can prove that the worst case occurs for $c = 1$ and $c_1 = n$ which gives us the worst-case time complexity $O(n \log n + l + T(n))$.

Memory complexity. All operations on input quads Q (URI resolution, sorting and removal of duplicates) can be executed in place. The only extra space is required for the disjoint-set representing the canonical URI mapping. Implementation as a tree with one node per URI requires $2l = O(l)$ nodes.

5. Fused Linked Data Quality Algorithm

Data quality can be perceived as "fitness of use" with respect to a particular task [14]. Traditionally, quality of data is assessed before the data fusion phase based on factors such as timeliness or accuracy [31, 34] relating to a single data source which prevents the fusion phase from consulting other sources. We propose a new approach that can leverage the context of multiple sources in quality assessment.

5.1. F-Quality

We introduce a new term *F-quality* for the special case of fused data quality in order to distinguish the quality of source data (expressing quality of data per se) and quality of fused data (expressing quality in the context of other – possibly conflicting – sources).

The purpose of F-quality is to (1) act as one of the deciding factors during conflict resolution, (2) help data

consumers decide which data is worth using, and (3) detect low quality data before further processing. These goals are further supported in our solution by accompanying fused data with provenance metadata indicating where a particular value came from and thus supporting verifiability of information.

Definition 9 (F-Quality score). *F-quality score* (or *F-quality* for short) is a number from interval $[0; 1]$ expressing the quality of a value after data fusion with respect to other conflicting values, content of the value, provenance of the value, and quality-related metadata.

The definition is intended for quality in the data integration context. Taking other conflicting values into consideration is what distinguishes it from how quality is regarded in this context, where only the rest of the factors are commonly used [34].

For the purposes of Linked Data integration, we can replace the term *value* in the definition with *triple* or *statement*. F-quality is meant to be a factor used for conflict resolution decisions and also as a lead for data consumer's trust decisions about different values. For this reason, it is necessary to assess F-quality at the statement or value level (rather than with lower granularity).

The range of F-quality needs to be from an ordered space so that values can be easily compared. We define it as the interval $[0; 1]$ where 0 means no confidence and 1 maximal confidence. This is in accordance with other related work where either the $[0; 1]$ interval is used [19, 32], or simple binary decisions are made [2]. Range $[-1; 1]$ has been proposed for the related concept of *trust* [11], which is different from quality, however.

As stated in Section 4.1, the resolution of RDF statements rather than property values is convenient for Linked Data. Since we focus on resolution of RDF triple objects, we do not further distinguish between F-quality of a resolved statement and its object.

5.2. F-Quality Assessment

We propose an F-quality assessment method based on three factors:

1. Quality of data sources.
2. Data conflicts.
3. Confirmation of values by multiple sources.

The exact steps of quality assessment may be customized for needs of each resolution function.

Quality of source data is modelled on the Named Graph level – Named Graphs are convenient for attaching metadata to a set of triples and are commonly used in quality assessment [2, 6, 19]. For a statement from a Named Graph (G, g), we interpret the quality score associated with (G, g) as the initial statement's quality.

Assessment of source Named Graphs' quality is *not* part of our assessment algorithm. Quality dimensions can be described by various indicators [31, 34] which need to be converted to a single quality score for our purposes. This task is beyond the scope of this paper and quality scores are simply an input value. A dedicated Quality Assessment module handles this task in ODCleanStore [16]; other solutions include [10, 11, 12, 19].

Methodology. We took the following steps when designing our assessment method.

First, we formulated objectives of F-quality, as outlined in Section 5.1: act as a deciding factor in conflict resolution, help data consumers in decisions, and detect low data quality. Another objective was to enrich existing quality assessment methods with the ability to leverage the conflict resolution context.

As the next step, we identified information that can be used during assessment. Authors in [2] classify quality indicators as based on information content itself, on metadata about the origin or on user ratings. The information content is represented by the conflicting statements, origin metadata and potential user ratings are modelled on the Named Graph level. They should be converted to a single score. Existing quality assessment tools can be utilized and the result quality score is given among metadata for the resolution function.

Next, we collected several real-world cases for data integration [16]. The objectives and use cases lead us to several requirements on the assessment method summarized in Section 5.3.

We came up with several functions that attempted to satisfy the requirements. We chose one of the functions, which emerged as a simple and natural solution of the requirements, implemented it and proved its feasibility on evaluation in ODCleanStore.

Terminology. Let us pick up on the formalism introduced in Section 4.1:

Definition 10 (F-Quality function q). Let q be a function $q : N \times \mathcal{P}(U) \times \mathcal{P}(Quads) \times \mathcal{P}(Quads) \rightarrow [0; 1]$.

We interpret the value $q(\mathbf{v}, S, CC, M)$ as the F-quality of value $\mathbf{v}$ (or resolved quad with $\mathbf{v}$ as its object, respectively). The value whose quality we calculate is represented by $\mathbf{v} \in N$. $S \subseteq U$ is the set of names of graphs which state this value (for deciding resolution functions[8]) or which $\mathbf{v}$ was derived from (for mediating

[8] See Definition 8.

functions[8]) such that $S \neq \emptyset$ and $S \subseteq graphs(CC)$. We refer to S as *source Named Graphs*. CC is the set of conflicting statements (*conflict cluster*) and M the set of metadata quads – this is the same as for a resolution function.

Definition 11 (Distance measure). Distance measure is a function $d : N \times N \rightarrow [0; 1]$ such that d is symmetric and $d(n, n) = 0$. Value $d(n_1, n_2)$ is interpreted as the distance or difference of values represented by RDF nodes n_1 and n_2 where 0 means most similar and 1 means completely different.

Definition 12 (Graph quality score). Let $s : U \rightarrow [0; 1]$ be a partial function which assigns quality score $s(g)$ to a Named Graph (G, g).

The value of $s(g)$ is supposed to be either precomputed and given in metadata M or it should be possible to calculate from M, therefore it doesn't need to be an explicit argument of the F-quality function q.

We will also use $\bar{s} : \mathcal{P}(U) \rightarrow [0; 1]$ to denote a function which aggregates scores of several Named Graphs based on $s(g)$.

5.3. Requirements on F-Quality Assessment

This section summarizes requirements on an F-quality function. We have identified mathematical properties that should be satisfied and three main quality factors: (1) quality scores of source Named Graphs, (2) value conflicts and (3) confirmation by multiple sources.

1. Function q doesn't depend on the order of elements in S or CC (symmetry).
2. Function q is increasing with $\bar{s}(S)$ (monotonicity).

$$\bar{s}(S_1) \geq \bar{s}(S_2) \Rightarrow q(\mathbf{v}, S_1, CC, M) \geq q(\mathbf{v}, S_2, CC, M)$$

3. $\bar{s}(S)$ is a function of $s(g_i)$, $g_i \in S$ and monotonous in its arguments ($\bar{s}$ monotonicity).

$$(S = \{g_1, \ldots, g_n\}, S' = \{g'_1, \ldots, g'_n\}, \forall i = 1, \ldots, n : s(g_i) \geq s(g'_i)) \Rightarrow \bar{s}(S) \geq \bar{s}(S')$$

4. Function q decreases with increasing distance from other conflicting values $objects(CC)$. Significance of such difference should be weighted by the respective graph quality scores.

$$(\forall v_i \in objects(CC) : d(\mathbf{v}, v_i) \geq d(\mathbf{v}', v_i)) \Rightarrow q(\mathbf{v}, S, CC, M) \leq q(\mathbf{v}', S, CC, M)$$

5. If $objects(CC) = \{\mathbf{v}\}$, then $q(\mathbf{v}, S, CC, M) \geq q(\mathbf{v}', S, CC, M)$ for all $v' \neq v$ (idempotency).
6. If there are no conflicts, i.e. $objects(CC) = \{\mathbf{v}\}$, then $q = \bar{s}(S)$.
7. Sources with zero quality score should not affect q.
8. If n sources with the maximum score claim a value completely different from $\mathbf{v}$, it should decrease value of q approximately n times.
9. If several sources agree on the same value $\mathbf{v}$, it increases the value of $q(\mathbf{v}, S, CC, M)$.

Requirements 1-5 constraint mathematical properties of q while requirements 6-9 were introduced so that q produces intuitive results. Our requirements overlap with framework proposed in [32] where a more detailed analysis of resulting properties of q can be found. See also related work in Section 3.3.

5.4. F-quality Assessment Algorithm Description

In this section, we describe an F-quality assessment method which complies with the requirements outlined in Section 5.3.

The assessment is called when a resolution function is executed. Therefore the details of the general assessment algorithm may be customized for the function – we use a different version for mediating and deciding resolution functions in ODCS-FusionTool, for instance.

The assessment algorithm is described in Listing 4, and details for each quality factor discussed below.

5.4.1. Factor 1: Source Quality

The first step is to use quality of sources (line 1). Sources are represented by a set S of URIs of Named Graphs value $\mathbf{v}$ originated from. Value of this parameter is resolution function dependent and supplied by the function. ODCS-FusionTool uses $S = \{g \mid (s, p, \mathbf{v}, g) \in CC\}$ for deciding resolution functions (graphs that actually contain $\mathbf{v}$), and all graphs $S = \{g \mid (s, p, o, g) \in CC\}$ for mediating functions. When using the AVG (average) resolution function, for instance, all values contributed to the average indeed and should be considered as sources.

Function $\bar{s}$ computing total score from $s(g_i)$, $g_i \in S$ is also resolution-specific. $\bar{s}$ must be monotonous and symmetric, which is satisfied, e.g., by minimum, maximum, or average. ODCS-FusionTool uses maximum for deciding functions and average for mediating functions.

We can continue with Example 3: In ODCleanStore, function $s(g_i)$ would return the value of `odcs:score` for g_i in metadata M. Therefore, the value of q after step one would be $q = 0.9$ (maximum for `dbpedia.org`

Listing 4: F-quality assessment algorithm

Input: Assessed value $\mathbf{v}$, source graph names (for $\mathbf{v}$) S, collection of conflicting statements $CC = \{(s_1, p_1, o_1, g_1), \ldots, (s_n, p_n, o_n, g_n)\}$, metadata M. The algorithm is also parametrized with a distance measure d, graph quality score functions s and $\bar{s}$, boolean constants $ConsiderConflicts$, $ConsiderConfirmation$ and a positive numeric constant $AgreeCoefficient$.

Output: F-quality of value $\mathbf{v}$.

1: $q \leftarrow \bar{s}(S)$ // Factor 1: Source quality
2: **if** $ConsiderConflicts$ **then** // Factor 2: Conflicting values
3: $conflictFactor \leftarrow 1 - \frac{\sum_{i=1}^{n} s(g_i) \cdot d(\mathbf{v}, o_i)}{\sum_{i=1}^{n} s(g_i)}$
4: $q \leftarrow q \cdot conflictFactor$
5: **end if**
6: **if** $ConsiderSupport$ **then** // Factor 3: Confirmation by multiple sources
7: $support = \{g_i \mid o_i = \mathbf{v}\}$
8: **if** $|support| > 0$ **then**
9: $supportFactor \leftarrow \frac{\left(\sum_{g_i \in support} s(g_i) - \max_{g_i \in support} s(g_i)\right)}{AgreeCoefficient}$
10: $supportFactor = \min(supportFactor, 1)$
11: $q \leftarrow q + (1 - q) \cdot supportFactor$
12: **end if**
13: **end if**
14: **return** q.

and `rdf.freebase.com`) for label $\mathbf{v}$ ="`Berlin`", and $q = 0.85$ (average) for average longitude $\mathbf{v} = 19.391$.

5.4.2. Factor 2: Conflicting Values

The second step (lines 2-5) looks at the actual data to be resolved and conflicts between them using a distance measure d. The more a value deviates from what others claim, the lower quality it has, taking source quality scores into consideration.

This step is executed optionally since different values may be acceptable in some cases, especially for many-valued properties such as `rdf:type`. There is a special user-provided resolution parameter which controls execution of this step in ODCS-FusionTool at the global or per-property level.

Conflict Factor. Line 3 calculates an average of value distances weighted by the respective quality scores. This can be interpreted as each source voting for value $\mathbf{v}$, its vote proportional to the similarity of o_i claimed by the source to $\mathbf{v}$ and the vote weighted by source quality score.

Multiplication in expression $s(g_i) \cdot d(\mathbf{v}, o_i)$ on line 3 and on line 4 could be replaced by any *and*-like operator. An *and*-like operator is a binary operator $\wedge$ such that $\wedge$ is monotonous, $0 \wedge x = x \wedge 0 = 0$ and $1 \wedge x = x \wedge 1 = x$. This is satisfied, e.g., by multiplication or minimum. We chose multiplication on line 3 because a weighted average is easily comprehensible and interpretable as "voting". Multiplication on line 4 was chosen over minimum to satisfy requirement 8 and for the reason demonstrated by the following example. If we have a source with quality score $s(g_i) = 0.5$ and $conflictFactor = 0.5$, then we can intuitively expect that the result quality should be strictly less than when $s(g_i) = 0.5$ and $conflictFactor = 1$ (or the other way round). Taking minimum wouldn't be able to distinguish these cases.

Distance Measure. Function d should be customized to the type of values it is applied to. Functions commonly employed in duplicate detection are best suited for literals, e.g., Levenshtein or Jaro-Winkler distance for strings [24, 30] or $\min(|v_1 - v_2|/max, 1)$ for ordinal values (max being a parameter) [24, 32]. A domain-specific measure can be utilized for URI resources, e.g., value based on color similarity can be returned for resources representing colors like `ex:black`, `ex:gray`, etc. ODCS-FusionTool uses Levenshtein for strings, $\min(|2(v_1 - v_2)/(v_1 + v_2)|, 1)$ for numbers (distance normalized by average), $\min(|v_1 - v_2|/max, 1)$ for dates and inequality indicator (0 for equal values, 1 for others) otherwise.

In our working example, quality q would decrease for label `"Berlin"` due to its difference from label `"Berlin (Germany)"`, and quality q for longitude would decrease due to the difference between average longitude and source longitude values.

5.4.3. Factor 3: Confirmation by Multiple Sources

This last step (lines 6-13) is motivated by requirement 9. If several sources agree exactly on value $\mathbf{v}$, than it should have a higher F-quality than any of the sources alone. Let us imagine three sources g_1, g_2, and g_3 each having source quality score 0.5 and each claiming the same value $\mathbf{v}$. Intuitively, one can trust the value $\mathbf{v}$ more than if there was only one source supporting it. In this example, the resulting F-quality would be $q = 0.5$ without this factor, while it yields $q = 0.75$ when $AgreeCoefficient$ is 2.

This step is parametrized with a constant $AgreeCoefficient \in \mathbb{R}^+$. The idea is that the F-quality of **v** should increase linearly with the sum of source quality scores $\sum_{g_i \in support} s(g_i)$ so that

- if $|support| = 1$, then q is unchanged,
- the maximum F-quality of 1 is reached when $AgreeCoefficient + 1$ sources agree on **v**; more formally when

$$AgreeCoefficient + \max_{g_i \in support} s(g_i) = \sum_{g_i \in support} s(g_i).$$

This step is also optional because it doesn't make sense for resolution functions such as Avg or Concat[9] where a value match would be coincidental and shouldn't increase the F-quality.

In our working example, value of q for label `"Berlin"` would increase because two sources agree on the value, but q for longitude wouldn't change because factor 3 is not used for the Avg resolution function.

5.5. Complexity & Relation to Resolution Functions

The time complexity of the F-quality assessment algorithm is linear in the number or quads in the given conflict cluster. This determines the time complexity of resolution functions which are summarized in Figure 3.

We assume that evaluation of $s(g)$ and $d(x, y)$ is $O(1)$, and $\bar{s}(S)$ is $O(|S|)$ (true in ODCleanStore). Factor 1 is executed in $O(|S|) = O(n)$ because $S \subseteq graphs(CC)$. Factor 2 has a loop with n iterations in constant time. Factor 3 can be implemented in one pass over the conflict factor. Therefore the total complexity is $O(n)$.

Figure 3 gives an overview of default resolution functions in ODCS-FusionTool, their time complexity and quality assessment settings. These results further assume that metadata M can be queried for quads with a known subject and predicate in $O(\log |M|)$, conflict cluster is sorted from the conflict resolution algorithm (Listing 1), and resolution functions are applicable to all values. The complexity is quadratic for resolution functions All, Best, TopN and Threshold because F-quality must be evaluated for all quads in a conflict cluster, each needing an $O(n)$ time. Other resolution functions evaluate quality assessment only for the one resolved quad they return.

[9] See Appendix A

5.6. Discussion

The proposed F-quality assessment algorithm covers the goals outlined above and gives results easy to interpret. The assessment algorithm is independent on the quality assessment of data sources which may be task-specific and is beyond the scope of this paper. A downside of the F-quality assessment is that considering context represented by the conflict cluster CC requires time $O(|CC|)$ which results in a quadratic complexity of some resolution functions.

The presented assessment algorithm provides a basic framework for quality assessment in a data integration environment, taking data conflicts into consideration. Information quality is strongly task-dependent [2], however. Data consumers have different requirements for a medical application and a simple music database. In addition, quality assessment needs to be tailored for the resolution function it is used with (Figure 3).

For this reason, quality calculation is not a part of the conflict resolution algorithm but it is a part of resolution functions. Custom resolution functions with custom quality formulas may be plugged-in and the formula suggested here may serve as a basis for custom-tailored methods.

6. Evaluation

We have implemented and evaluated the proposed data fusion and quality assessment algorithms in the ODCleanStore framework and also in a standalone tool ODCS-FusionTool. The experience from more than a year of practical use has proven their feasibility and usefulness. In this section, we present our evaluation on publicly available Linked Data used in the OAEI evaluation campaign. We investigate both performance and effectiveness of the presented algorithms.

For the evaluation, we chose the dataset used for the Instance Matching challenge of OAEI 2011[10] including data about locations from The New York Times (NYT), Freebase, DBPedia, and Geonames datasets because NYT contains high-quality curated links among these datasets and therefore the quality of links doesn't impair the results. The data were supplemented with ontology mappings available for Geonames and mappings handcrafted for the most frequent properties and classes in the datasets. The dataset contained information about 3840 locations from NYT described by 2 million triples.

The prepared data were integrated with three resolution functions: All, Any and Best. Figure 4 summarizes dataset statistics before and after integration.

[10] http://oaei.ontologymatching.org/2011/instance/

Type	*Consider-Conflicts*	*Consider-Support*	Resolution function	Time complexity
Deciding	yes	yes	`All, Best, TopN, Threshold`	$O(n^2)$
			`Any, BestSource, Filter, Longest, Max, Min, None, Shortest, Vote, WeightedVote`	$O(n)$
			`MaxSourceMetadata, MinSourceMetadata`	$O(n \log \|M\|)$
Mediating	yes	no	`Avg, Median`	$O(n)$
Mediating	no	no	`Concat, Sum`	$O(n)$

Figure 3: Complexity and F-quality settings of resolution functions in ODCleanStore. See Appendix A for description of functions.

	NYT	DBPedia	Freebase	Geonames	All data	**All**	**Best/Any**
Total triples	117,623	260,426	1,567,454	55,535	2,005,058	1,461,449	461,449
`owl:sameAs` links	17,572	41,702	0	26	67,340	67,340	67,340
All subjects	5,620	2,100	1,917	3,544	11,481	6,624	6,624
Unique subjects	3,832	2,083	1,917	3,544	6,624	6,624	6,624
Conflict clusters	55,894	168,188	268,289	43,296	456,957	461,449	461,449
Average cluster size	2.10	1.55	5.84	1.28	4.39	4.11	1
Unique predicates	36	2,870	1,323	33	4,221	4,221	4,221

Figure 4: Evaluation dataset statistics before and after integration. "All data" include additional `owl:sameAs` mappings.

6.1. Completeness, Conciseness, Consistency

Increasing *completeness*, *conciseness* and *correctness* are three broad goals of data integration [9]. Correctness expresses how much the data conform to the real world while completeness and conciseness are in a way analogous to recall and precision in information retrieval.

Completeness measures the amount of data relative to all available data in the given domain and is achieved by adding more data. We distinguish extensional and intensional completeness defined by the following formulas:

$$\textit{ext. completeness} = \frac{|\ \textit{unique objects in dataset}\ |}{|\textit{unique objects in universe}|}$$

$$\textit{int. completeness} = \frac{|\textit{unique attributes in dataset}|}{|\ \textit{unique available attributes}\ |}$$

Conciseness measures the uniqueness of object representations. Conciseness is increased when redundant data are removed. Again, we recognize extensional and intensional conciseness:

$$\textit{ext. conciseness} = \frac{|\textit{unique objects in dataset}|}{|\textit{all objects in dataset}|}$$

$$\textit{int. conciseness} = \frac{|\textit{unique attributes in dataset}|}{|\textit{all attributes in dataset}|}$$

Consistency, as defined in [19], is related to correctness but easier to measure. It measures conflicts in the datasets regardless of correct real world values.

$$\textit{consistency} = \frac{|\textit{conflict clusters without conflicts}|}{|\textit{all conflict clusters}|}$$

For the purposes of our evaluation, we identify the number of objects in the universe with the number of unique entities in the source datasets combined; we also do the same with attributes. Consistency is measured as the ratio of conflict clusters without conflict (i.e. with a single unique object) to the total number of conflict clusters *excluding* clusters for manyvalued properties. These properties have cardinality larger than one and their different values doesn't affect consistency.

Results. Measurements for each of the source datasets and all data combined are listed in Figure 5. The listed numbers are relative to the result integrated using the Best resolution function – values of completeness, conciseness and consistency are therefore 100% after integration for all three resolution functions except for consistency of 78.9% measured on data resolved by All.

We can see that completeness is raised by simply putting data together. However, this results in decrease of conciseness and consistency which must be solved by

	NYT	DBPedia	Freebase	Geonames	All data	**ALL**	**BEST/ANY**
Ext. completeness	57.9%	31.5%	29.0%	53.5%	100.0%	100.0%	100.0%
Int. completeness	0.9%	68.0%	31.3%	0.8%	100.0%	100.0%	100.0%
Ext. conciseness	68.2%	99.2%	100.0%	100.0%	57.7%	100.0%	100.0%
Int. conciseness	100.0%	100.0%	100.0%	100.0%	97.3%	100.0%	100.0%
Consistency	90.9%	88.5%	73.4%	97.0%	78.7%	78.9%	100.0%

Figure 5: Completeness, Conciseness and Consistency of the evaluated data sets relative to the result of integration using BEST.

	ALL	BEST	ANY
Initialization	0:01.0	0:00.9	0:01.0
Triple loading	8:45.8	8:36.8	8:32.7
Conflict Resolution	0:44.2	0:38.7	0:10.8
Total	9:52.2	9:37.1	9:03.7
CR throughput (triple/s)	45,360	51,810	185,650
CR throughput (CCs/s)	10,334	11,807	42,310
Total throughput (triple/s)	771	791	863

Figure 6: Run times of on sample datasets (min:sec). "CCs" stands for conflict clusters.

the data fusion process. ODCS-FusionTool raised extensional conciseness by 73% and intentional conciseness by 3%. Improvement in consistency depends on configuration of conflict resolution – while returning all conflicting values has a minimal impact on consistency, choosing a single best value improves it by 27% in our evaluation.

A similar experiment on a dataset of Brazilian municipalities was conducted for Sieve in [19]. Their results apply to ODCS-FusionTool as well because it can be configured to produce the same results on the dataset.

6.2. Performance

The runtime of each integration run was measured on a server with four 3.00GHz CPUs and 16 GB of memory. Figure 6 lists the results. The run times are median values over three runs.

Several facts can be observed from the results. One is that the initialization phase, which includes resolution of canonical URIs, is very fast. The resolution would scale well even for a large number of `owl:sameAs` links thanks to its nearly linear time complexity.

Most time is consumed by triple loading. This issue deserve more attention in future work. Note that this problem is visible only when processing data in large batches. This is not the case in ODCleanStore where conflicts are resolved on demand for a single query.

Run times of the actual data fusion/conflict resolution part differ for ALL and BEST resolution functions, and for resolution function ANY. This is according to expectations – the former functions have quadratic time complexity in the size of a conflict cluster while the latter is linear (see Figure 3). The throughput is satisfactory and conflict resolution scales linearly for typical data – even for large number of triples, the size of conflict clusters is typically limited by a reasonable constant. The average size of a conflict cluster was 4.39 on the test data, with maximum of 279 triples in a single cluster.

6.2.1. Quality

In this section, we look closer at F-quality values produced by our algorithm. We selected several properties from resource description of the city of Berlin in the evaluation dataset. F-quality calculation uses metadata about quality of sources. We added sample values expressing preference of data from DBPedia (with score 0.9) over other sources (with score 0.8):

```
<\protect\vrule width0pt\protect\href{http://dbpedia.or
<\protect\vrule width0pt\protect\href{http://rdf.freeba
<\protect\vrule width0pt\protect\href{http://sws.geonam
<\protect\vrule width0pt\protect\href{http://data.nytim
<\protect\vrule width0pt\protect\href{http://example.co
```

In order to demonstrate behavior in occurrence of errors, we also added an artificial source "Err" claiming an incorrect value of geographical latitude of Berlin. Figure 7 lists values of the selected properties together with provenance information and F-quality produced by ODCS-FusionTool for the sample data. The resolution function used is ALL.

If we look at F-quality of labels, we see `"Berlin"` is clearly the best value. This is expected as three sources agree on it. F-quality of other labels depends mostly on their similarity to `"Berlin"`.

The errorneous value of latitude has been clearly determined as low-quality because it is in conflict with other sources. Latitude from the preferred source (DBPedia) is ranked highest. The remaining values are

Property	Value	F-Quality	Sources
rdfs:label	"Berlin"	0.75992	DBPedia, Freebase, Geonames
	"City_of_Berlin"	0.27447	Freebase
	"Berlin (Germany)"	0.22126	NYT
geo:lat	"52.5006"	0.72418	DBPedia
	"52.5167"	0.64381	NYT
	"52.5233"	0.64380	Freebase
	"52.52437"	0.64380	Geonames
	"13.4126"	0.15610	Err
geo:long	"13.3989"	0.89957	DBPedia
	"13.4"	0.79965	NYT
	"13.41053"	0.79963	Geonames
	"13.4127"	0.79956	Freebase
dbprop:web	<http://www.berlin.de...php>	0.37739	DBPedia,Freebase
	<http://berlin.unlike.net/>	0.11793	DBPedia
	<http://www.berlin.de>	0.09275	Freebase
	...		
rdf:type	schema:City	0.92000	DBPedia,Freebase
	schema:Place	0.90000	DBPedia
	geonames:Feature	0.80000	Geonames
	...		

Figure 7: Sample F-quality values for integrated resource description of Berlin

ranked according to how close they are to the (weighted) average of all values. The same holds for longitude.

The website suggested by two sources is ranked best while remaining values are ranked according to quality score of their source.

Values of `rdf:type` are special in that the cardinality parameter was set to `MANYVALUED` for them. This corresponds to $ConsiderConflicts = false$ in the F-quality algorithm, therefore conflicts between values do not decrease the quality. One can see that F-quality then depends only on score of the underlying source and support by multiple sources.

It is important to realize that the relative F-quality for each value is more important than absolute values. The purpose of F-quality was established to be (1) a deciding factor for conflict resolution, (2) a decision support for users and (3) an indicator of low-quality data. This goal was indeed achieved on the sample data. Using the `Best` resolution function would really return the most reasonable value for each property, a user would be given a good lead on which website is worth visiting and the errorneous value of latitude can be detected.

7. Specifics of Linked Data Integration

We conducted a thorough comparison of Linked Data integration compared to relational databases as part of our work on ODCS-FusionTool. This section summarizes its basic specifics and open challenges.

Resource descriptions. Data fusion deals with *records* [3]. While this term naturally matches tuples in relational databases, its equivalent in Linked Data is ambiguous. The most common approach is with *resource descriptions* composed of RDF triples having the given RDF resource as its subject. It is sufficient in most scenarios and enables straightforward application of traditional database data fusion techniques. Other options are possible, such as the inclusion of adjacent blank nodes [27]. Including triples having the resource of interest in place of object is also an option.

Blank nodes. Blank nodes (RDF resources without an identifier) introduce additional complexity in Linked Data integration across several steps. One problem is with query "round-tripping" which affects the data retrieval step. If a result of a SPARQL query contains a blank node `_:b1` in an object and we want to retrieve the resource description of `_:b1`, we cannot reference it in a subsequent query because `_:b1` has no longer any relation to the source graph [18].

The lack of global identifiers makes it impossible for duplicate detection to produce mappings between them.

The data fusion step can interpret blank nodes as either autonomous entities, or structured properties. In the latter case, it is possible to treat the structured property components independently or together.

Finally, our F-quality assessment algorithm uses a distance measure to compare values – comparing blank nodes as structured attributes is computationally difficult [5].

Schemata. The equivalent of schema in RDF is an ontology, typically expressed in OWL. While ontology can define restrictions, they are often not applied strictly [13]. Schemata are rather loosely used, and properties and classes need not be defined explicitly. This makes automation difficult (e.g., automatic generation of conflict resolution settings [4]). Another implication is that satisfaction of integrity constraints cannot be reliably used to detect low quality data. On the other hand, this approach is very flexible and target schema can evolve more easily. Schema mappings can also be richer (e.g., subproperty/subclasss taxonomies).

Query Execution. Implementations of the fusion step in relational databases are built on non-trivial join- and union-based operations [3] in order to align schema and join data from different sources. RDF simplifies this aspect to simply selecting relevant triples. In addition, we have the power of SPARQL and inferencing at hand.

Nulls. The special value `NULL` has no equivalent in RDF other than simply a missing value. This makes it difficult to distinguish uncertainty and contradiction – e.g., when one source contains a book's reviewer while other doesn't, the reviewer may be unknown (uncertainty) or the book wasn't reviewed (contradiction).

8. Conclusion

A wider adoption of Linked Data needs a critical amount of data, experience, and tools. This paper contributes to the arsenal of available tools with ODCS-FusionTool, a data fusion and conflict resolution tool with quality assessment and provenance tracking. We also contribute to the data integration experience with a summary of Linked Data integration specifics and an overview of possible conflict resolution functions (Appendix A).

Our implementation is based on new well-defined algorithms. One algorithm fuses data with resolution of identifier, schema, and data conflicts. The other one computes fused data quality for which we introduce a new term *F-quality* – it is novel in quality assessment of fused data rather than source data, and leveraging the context of multiple (possibly conflicting) sources. The combination of the two algorithms is mutually beneficial and enables data fusion to make quality-based decisions even in the absence of quality-related metadata.

ODCS-FusionTool has been evaluated on real-world Open Linked Data. We achieved a distinct improvement of conciseness and consistency and gained valuable observation for future research. Usefulness and feasibility has also been proven by more than a year of active usage in relation to the OpenData.cz initiative[11] on processing of Czech public contracts data and EU procurement notices.[12]

8.1. Future work

One area for further research are blank nodes. The proposed data fusion algorithm uses triples sharing the same subject as resource descriptions. An open question is how to effectively leverage inclusion of blank nodes in resource descriptions and decide between interpreting them as entities or structured attributes.

An interesting option is integration of conflict resolution with the SPARQL query language like other systems extend SQL [1, 22, 33]. Currently, data need to be loaded beforehand and conflict resolution settings have no effect when querying the underlying RDF store.

Our experiments show data retrieval as the performance bottleneck. SPARQL queries executed on the underlying RDF store constituted up to 90% of execution time. More efficient access to data, caching and parallelization should be investigated. Our algorithm is parallelizable in that once canonical URIs are resolved, resource descriptions are resolved independently.

References

References

[1] A. Bilke, J. Bleiholder, F. Naumann, et al., Automatic data fusion with HumMer, in: Proceedings of the 31st International Conference on Very Large Data Bases, VLDB Endowment, 2005.

[2] C. Bizer, R. Cyganiak, Quality-driven information filtering using the WIQA policy framework, Web Semantics: Science, Services and Agents on the World Wide Web 7 (1) (2009) 1–10.

[3] J. Bleiholder, Data Fusion and Conflict Resolution in Integrated Information Systems, Ph.D. thesis, Universitt Potsdam (2010).

[4] V. Bryl, C. Bizer, Learning conflict resolution strategies for cross-language wikipedia data fusion, in: 4th Joint WICOW/AIRWeb Workshop on Web Quality Workshop (WebQuality) – WWW 2014, 2014.

[11] `http://opendata.cz`

[12] `http://isvzus.cz`, `ted.europa.eu`

[5] J. J. Carroll, Signing rdf graphs, in: The Semantic Web – ISWC 2003, Springer, 2003, pp. 369–384.
[6] J. J. Carroll, C. Bizer, P. Hayes, P. Stickler, Named Graphs, Provenance and Trust, in: Proceedings of the 14th international conference on World Wide Web, ACM, 2005.
[7] S. Chawathe, H. Garcia-Molina, J. Hammer, et al., The TSIMMIS Project: Integration of Heterogenous Information Sources, in: Proceedings of IPSJ, 1994.
[8] U. Dayal, Processing Queries Over Generalization Hierarchies in a Multidatabase System, in: Proceedings of the 9th International Conference on Very Large Data Bases, VLDB '83, Morgan Kaufmann Publishers Inc., 1983.
[9] X. L. Dong, F. Naumann, Data Fusion: Resolving Data Conflicts for Integration, Proceedings of the VLDB Endowment 2 (2) (2009) 1654–1655.
[10] A. Flemming, Quality Characteristics of Linked Data Publishing Datasources, Master's thesis, Humboldt-Universität zu Berlin (2010).
[11] O. Hartig, Trustworthiness of Data on the Web, in: Proceedings of the STI Berlin & CSW PhD Workshop, Citeseer, 2008.
[12] O. Hartig, J. Zhao, Using Web Data Provenance for Quality Assessment, in: Proceedings of the International Workshop on Semantic Web and Provenance Management, Washington DC, USA, 2009.
[13] A. Hogan, A. Harth, A. Passant, S. Decker, A. Polleres, Weaving the Pedantic Web, in: 3rd International Workshop on Linked Data on the Web, LDOW 2010, CEUR, 2010.
[14] J. M. Juran, A. B. Godfrey, Juran's Quality Handbook, 5th ed., McGraw-Hill, New York, 1998.
[15] T. Knap, J. Michelfeit, J. Daniel, P. Jerman, D. Rychnovský, T. Soukup, M. Nečaský, ODCleanStore: A Framework for Managing and Providing Integrated Linked Data on the Web, in: Web Information Systems Engineering – WISE 2012, Springer, 2012, pp. 815–816.
[16] T. Knap, J. Michelfeit, M. Nečaský, Linked Open Data Aggregation: Conflict Resolution and Aggregate Quality, in: Computer Software and Applications Conference Workshops (COMPSACW), 2012 IEEE 36th Annual, IEEE, 2012.
[17] E.-P. Lim, R. H. Chiang, A Global Object Model for Accommodating Instance Heterogeneities, in: Conceptual Modeling – ER 1998, Springer, 1998, pp. 435–448.
[18] A. Mallea, M. Arenas, A. Hogan, A. Polleres, On blank nodes, in: The Semantic Web – ISWC 2011, Springer, 2011, pp. 421–437.
[19] P. N. Mendes, H. Mühleisen, C. Bizer, Sieve: Linked Data Quality Assessment and Fusion, in: Proceedings of the 2012 Joint EDBT/ICDT Workshops, ACM, 2012.
[20] A. Motro, P. Anokhin, Fusionplex: resolution of data inconsistencies in the integration of heterogeneous information sources, Information fusion 7 (2) (2006) 176–196.
[21] N. W. Paton, K. Christodoulou, A. A. Fernandes, B. Parsia, C. Hedeler, Pay-as-you-go Data Integration for Linked Data: opportunities, challenges and architectures, in: Proceedings of the 4th International Workshop on Semantic Web Information Management, ACM, 2012.
[22] K.-U. Sattler, S. Conrad, G. Saake, Adding Conflict Resolution Features to a Query Language for Database Federations, in: Proceedings of the 3rd Workshop on Engineering Federated Information Systems, EFIS 2000, 2000.
[23] M. Scannapieco, A. Virgillito, C. Marchetti, et al., The DaQuinCIS Architecture: A Platform for Exchanging and Improving Data Quality in Cooperative Information Systems, Information Systems 29 (7) (2004) 551–582.
[24] E. Schallehn, K.-U. Sattler, G. Saake, Efficient similarity-based operations for data integration, Data & Knowledge Engineering 48 (3) (2004) 361–387.
[25] S. Schenk, S. Staab, Networked Graphs: A Declarative Mechanism for SPARQL Rules, SPARQL Views and RDF Data Integration on the Web, in: Proceedings of the 17th international conference on World Wide Web, ACM, 2008.
[26] A. Schultz, A. Matteini, R. Isele, C. Bizer, C. Becker, LDIF – Linked Data Integration Framework, in: 2nd International Workshop on Consuming Linked Data, Bonn, Germany, 2011.
[27] P. Stickler, CBD - Concise Bounded Description, Submission, W3C, `http://www.w3.org/Submission/CBD/` (2005).
[28] V. S. Subrahmanian, S. Adali, A. Brink, et al., HERMES: A Heterogeneous Reasoning and Mediator System (1995).
[29] R. E. Tarjan, J. Van Leeuwen, Worst-Case Analysis of Set Union Algorithms, Journal of the ACM (JACM) 31 (2) (1984) 245–281.
[30] J. Volz, C. Bizer, M. Gaedke, G. Kobilarov, Silk – A Link Discovery Framework for the Web of Data, in: Proceedings of the 2nd Linked Data on the Web Workshop, Citeseer, 2009.
[31] R. Y. Wang, D. M. Strong, Beyond Accuracy: What Data Quality Means to Data Consumers, Journal of Management Information Systems (1996) 5–33.
[32] R. R. Yager, A framework for multi-source data fusion, Information Sciences 163 (1) (2004) 175–200.
[33] L. L. Yan, M. T. Ozsu, Conflict Tolerant Queries In AURORA, in: Proceedings of IFCIS International Conference on Cooperative Information Systems, IEEE, 1999.
[34] A. Zaveri, A. Rula, A. Maurino, R. Pietrobon, J. Lehmann, S. Auer, Quality Assessment Methodologies for Linked Open Data, Submitted to The Semantic Web Journal.

Appendix A. List of Conflict Resolution Functions

A number of functions for resolution of conflicting values has been proposed in the literature [1, 3, 15, 20, 28, 33]. This appendix contains a comprehensive overview of functions relevant for Linked Data proposed in the literature with a few new additions. Some systems also allow custom user-defined resolution functions [22].

All. Returns all values.

Any. Returns an arbitrary (non-`NULL`) value.

First, Last. Returns the first or the last (non-`NULL`) value, respectively. Requires ordering of the values on input.

Random. Returns a random (non-`NULL`) value. The chosen value differs among calls on the same input.

Certain. If input values contain only one distinct (non-`NULL`) value, returns it. Otherwise returns `NULL` or empty output (depending on the underlying data model).

Best. Returns the value with the highest data quality value. The quality measure is application-specific.

TopN. Returns n best values (see Best). n is a parameter.

Threshold. Returns values with data quality higher then a given threshold. The threshold is given as a parameter.

BestSource. Returns a value from the most preferred source. The preference of source may be explicit (given preferred order of sources) or based on an underlying data quality model.

MaxSourceMetadata. Returns a value from the source with a maximal source metadata value. The metadata value may be, e.g., timestamp of the source, access cost or a data quality indicator. The used type of source metadata is either given as a parameter or fixed.

MinSourceMetadata. Returns a value from the source with the minimal source metadata value (see MaxSourceMetadata).

Latest. Returns the most recent (non-NULL) value. Recency may be available from another attribute, value/entity metadata or source metadata (the last case is a special case of MaxSourceMetadata).

ChooseSource. Returns a value originating from the source given as a parameter.

Vote. Returns the most-frequently occurring (non-NULL) value. Different strategies may be employed in case of tie, e.g., choosing the first or a random value.

WeightedVote. Same as Vote but each occurrence of a value is weighted by the quality of its source.

Longest, Shortest. Returns the longest/shortest (non-NULL) value.

Max, Min. Returns the maximal/minimal (non-NULL) value according to an ordering of input values.

Filter. Returns values within a given range. The minimum and/or maximum are given as parameters.

MostGeneral. Returns the most general value according to a taxonomy or ontology.

MostSpecific. Returns the most specific value, according to a taxonomy or ontology (if the values are on a common path in the taxonomy).

Concat. Returns a concatenation of all values. The separator of values may be given as a parameter. Annotations such as source identifiers may be added to the result.

Constant. Returns a constant value. The constant may be given as a parameter or be fixed (e.g. NULL).

CommonBeginning. Returns the common substring at the beggining of conflicting values.

CommonEnding. Returns the common substring at the end of conflicting values.

TokenUnion. Tokenizes the conflicting values and returns the union of the tokens.

TokenIntersection. Tokenizes the conflicting values and returns the intersection of the tokens.

Avg. Returns the average of all (non-NULL) input values.

Median. Returns the median of all (non-NULL) input values.

Sum. Returns the sum of all (non-NULL) input values.

Count. Returns the number of distinct (non-NULL) values.

Variance, StdDev. Returns the variance or standard deviation of values, respectively.

ChooseCorresponding. Returns the value that belongs to an entity (resource) whose value has already been chosen for an attribute A, where A is given as a parameter.

ChooseDepending. Returns the value that belongs to an entity (resource) which has a value v of an attribute A, where v and A are given as parameters.

MostComplete. Returns the (non-NULL) value from the source having fewest NULLs for the respective attribute across all entities.

MostDistinguishing. Returns the most distinguishing value among all present values for the respective attribute.

Lookup. Returns a value by doing a lookup into the source given as a parameter, using the input values.

MostActive. Returns the most often accessed or used value.

GlobalVote. Returns the most-frequently occurring (non-NULL) value for the respective attribute among all entities in the data source.